

ABSTRACT OF THE DISCLOSURE

[1069] We explore techniques for designing nonblocking algorithms that do not require advance knowledge of the number of processes that participate, whose time complexity and space consumption both adapt to various measures, rather than being based on predefined worst-case scenarios, and that cannot be prevented from future memory reclamation by process failures. These techniques can be implemented using widely available hardware synchronization primitives. We present our techniques in the context of solutions to the well-known *Collect* problem. We also explain how our techniques can be exploited to achieve other results with similar properties; these include long-lived renaming and dynamic memory management for nonblocking data structures.